Attorney's Docket No.: 10559/552001 / P12572

APPLICATION

FOR

UNITED STATES LETTERS PATENT

TITLE: SIMULTANEOUS BI-DIRECTIONAL CHANNEL SEPARATION

APPLICANT: STEPHEN R. MOONEY, AARON K. MARTIN, MATTHEW B. HAYCOCK, BRYAN K. CASPER, SHEKHAR Y. BORKAR, JOSEPH T. KENNEDY AND JAMES E. JAUSSI

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL584811582US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the U.S. Patent and Trademark Office, P.O Box 2327, Arlington, VA 22202.

November 28, 2001
Date of Deposit

Signature

~~Roxanne Ippolito~~ Gabe Lewis
Typed or Printed Name of Person Signing Certificate

SIMULTANEOUS BI-DIRECTIONAL CHANNEL SEPARATION

TECHNICAL FIELD

This invention relates to data transmission along a simultaneous bi-directional bus.

BACKGROUND

Busses are used to transmit data between various components in a computer or other device. For example, busses may be used to transmit data between a central processing unit, chipsets, memory, and peripheral components. Data transmission along busses occurs over discrete periods of time that are allocated for use by the components connected to the bus.

Typically, a pin is associated with each data line of a bus. Thus, as the width of busses increases, the number of pins needed by a component to be connected to a bus increases. Large numbers of pins hinder miniaturization of the components connected to the bus. Moreover, as the number of components connected to a bus increases, the timing and allocation of transmissions along the bus becomes increasingly complex and may ultimately slow down the operational speed of the device.

DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of a simultaneous bi-directional bus and a pair of devices communicating along the bus.

FIG. 2 is a graph of an exemplary frequency division on a simultaneous bi-directional bus.

FIG. 3 is a graph of an exemplary output signals and the combined signal on a simultaneous bi-directional bus.

FIG. 4 is a block diagram of another simultaneous bi-directional bus and a pair of devices communicating along the bus.

FIG. 5 is an exemplary process flow for arbitrating between two devices on a simultaneous bi-directional bus.

FIG. 6 is a block diagram of another simultaneous bi-directional bus and a pair of devices communicating along the bus.

Like reference symbols in the drawings indicate like elements.

DETAILED DESCRIPTION

Referring to FIG. 1, a first device 10 and a second device 20 communicate along a simultaneous bi-directional bus 30. Simultaneous bi-directional bus 30 carries a frequency division multiplexed signal between devices 10 and 20, which allows devices 10 and 20 to communicate simultaneously along a single conductor.

Devices 10 and 20 include, respectively, functional portions 50 and 60, which may be, for example, processors, memory storage units, control units, logic units, computational units, communication units, or other known units. Thus a functional portion 50 or 60 may be, for example, a microprocessor, chipsets (for example memory controllers, graphics controllers, and I/O controllers), a memory device, or an encoder or decoder for further data communication with a peripheral device. The devices 10 and 20 may each correspond to a single chip.

Functional portions 50 and 60 communicate over simultaneous bi-directional bus 30 using a transmitter 100 and receiver 150 of device 10, and a transmitter 200 and a receiver 250 of device 20. Transmitter 100 includes a high

pass encoder 110, a buffering transmitter 120, and a high pass filter 130 in series, while transmitter 200 includes a low pass encoder 210, a buffering transmitter 220, and a low pass filter 230 in series. Receiver 150 includes a low pass filter 160, a buffering receiver 170, and a low pass decoder 180 in series, while receiver 250 includes a high pass filter 260, a buffering receiver 270, and a high pass decoder 280 in series.

The cutoff frequencies and encoding procedures of transmitters 100, 200 and receivers 150, 250 are selected so that data from transmitter 100 is decoded at receiver 250 and data from transmitter 200 is decoded at receiver 150. For example, high pass encoder 110 uses combinational logic to encode outbound messages from functional portion 50 to produce high frequency encoded message signals. To do this, high pass encoder 110 encodes outbound messages with a maximum run length and symbol shapes with selected spectral content in order to minimize the low frequency content of the high frequency encoded message signals. As used herein, run length represents the number of consecutive symbols with the same value.

The encoded message signals are augmented by transmitter 120, which may be implemented as, for example, a CMOS current mode transmitter or a voltage mode transmitter. GaAs- and SiGe-based transmitters may also be used. High pass filter 130 filters off any spurious harmonics or low frequency noise in the high frequency encoded message signals to ensure that the encoded message signals do not interfere with data transmission between transmitter 200 and receiver 150.

The high frequency encoded message signals from transmitter 100 are received by low pass filter 160 and high pass filter 260. Low pass filter 160 filters out the high frequency encoded message signals to prevent propagation into

receiver 150. By contrast, high pass filter 260 passes the high frequency encoded message signals to buffering receiver 270, which boosts them and passes them along for decoding by high pass decoder 280. High pass decoder 280 uses complementary combinational logic to decode the high frequency encoded message signals into inbound messages for receipt by functional portion 60.

Communications from functional portion 60 to functional portion 50 are transmitted simultaneously along the same bi-directional bus 30 at low frequencies. Low pass encoder 210 uses combinational logic to encode outbound messages from functional portion 60 to produce low frequency encoded message signals. The encoded message signals are augmented by transmitter 220, which may be implemented as a CMOS current mode transmitter or a voltage mode transmitter. GaAs- and SiGe-based transmitters may also be used. Low pass filter 230 removes any spurious harmonics or high frequency noise in the low frequency encoded message signals to ensure that the low frequency encoded message signals do not interfere with data transmission between transmitter 100 and receiver 250.

The low frequency encoded message signals from transmitter 200 are received by low pass filter 160 and high pass filter 260. High pass filter 260 filters out the low frequency encoded message signals to prevent propagation into receiver 250. By contrast, low pass filter 160 passes the low frequency encoded message signals to buffering receiver 170, which boosts them and passes them along for decoding by low pass decoder 180. Low pass decoder 180 uses complementary combinational logic to decode the low frequency encoded message signals into inbound messages for receipt by functional portion 50.

An exemplary division of the frequency spectrum on bi-directional bus 30 is illustrated in FIG. 2. The spectral content of high frequency encoded message signals is within band 300, while the spectral content of low frequency encoded message signals is within band 310. Since bands 300 and 310 occupy different portions of the frequency spectrum, simultaneous data transmission along a single line can occur.

In FIG. 2, the bandwidth of band 300 is larger than that of band 310. This asymmetry is not necessary although it is advantageous when the majority of communication along bus 30 is in one direction. For example, a non-volatile memory may transmit significantly more data to a processor than it receives from the processor. In this case, band 300 is dedicated to transmissions from the non-volatile memory, and band 310 is dedicated to transmissions from the processor.

Exemplary time traces of a high frequency encoded message signal 320, a low frequency encoded message signal 330, and a combined signal 340 are illustrated in FIG. 3. Combined signal 340 is the signal on bi-directional bus 30 and results from the superposition of signals 320 and 330. High frequency filter 260 passes frequency encoded message signal 320 and low frequency filter 160 passes low frequency encoded message signal 330. Message signals 320, 330 can be, for example, control data, memory request data, read data, or write data.

High frequency encoded message signal 320 has a trapezoidal symbol shape and low frequency encoded message signal 330 has an exponential symbol shape. Other symbol shapes such as, for example, Gaussian, wavelet, and raised cosine pulse symbols may also be used.

Referring back to FIG. 1, in the illustrated implementation, the passbands and cutoff frequencies of filters 130, 160, 230, 260 and the combinational logic used by

encoders 110, 210 are predetermined. In other words, devices 10 and 20 are manufactured to transmit and receive data in predetermined spectral windows along simultaneous bi-directional bus 30.

In some implementations, each device 10, 20 may include transmitters and receivers with the particular frequencies used by each device being selected or tuned on an activity-by-activity basis. Thus, for example, a processor may use the larger band 300 when writing to a volatile memory, while the volatile memory may use the larger band 300 when being read by the processor.

For example, in the implementation shown in FIG. 4, the passbands and cutoff frequencies of filters 131, 161, 231, 261 and the combinational logic used by encoders 111, 211 are arbitrated between devices 11 and 21. Devices 10 and 20 each include a respective arbitration portion 191, 291 that arbitrates over an arbitration bus 301 connecting devices 11 and 21. Arbitration portions 191, 291 set the transmission and reception characteristics of the respective of transmitters 101, 210 and receivers 151, 251. To this end, filters 131, 161, 231, and 261 are programmable and encoders 111, 181, 211, and 281 have combinational logic tables that include inputs originating from arbitration portions 191, 291, or otherwise respond to signals from arbitration portions 191, 291.

FIG. 5 illustrates a process flow 500 for arbitrating between two devices. Initially, the first and second devices are powered on (step 510). Powering on triggers an arbitration process in which the first and second devices arbitrate to identify which will act as a high pass device and which will act as a low pass device. Determining a high pass device and a low pass device includes, for example, adapting

bandwidths to the anticipated need for data transmission between functional portions 50 and 60, establishing encoding protocols, and determining cutoff frequencies for the selected bandwidths (step 520). Next the identified high pass device sets a high pass transmitter and a low pass receiver (step 530), and the identified low pass device sets a low pass transmitter and a high pass receiver (step 540). Setting high and low pass transmitters and receivers includes setting the hardware components of the transmitters and receivers to function with the arbitrated characteristics. Finally, the high pass device and low pass device transmit data along a simultaneous bi-directional bus (step 550).

The arbitration of process flow 500 is performed after each system power up. Arbitration can be performed more or less frequently. For example, the results of a single arbitration can be stored in a non-volatile memory and revisited after power up, or arbitration can be performed on a periodic basis to ensure rate adaptation, i.e., to accommodate changes in the type of communication along a simultaneous bi-directional bus.

Referring to FIG. 6, in another implementation, the passbands and cutoff frequencies of filters 132, 162, 232, 262 and the format of the combinational logic used by encoders 112, 212 are selected by the user. Devices 10 and 20 each include a respective selection portion 192, 292 that relies upon a user's selection made at a selection device 401 to set the transmission and reception characteristics of the respective transmitters 102, 202 and receivers 152, 252. Filters 132, 162, 232, 262 are programmable and encoders 112, 212, 182, 282 have combinational logic tables that include inputs originating from selection portions 192, 292, or otherwise respond to signals from selection portions 192, 292.

Selection device 401 is illustrated as a single-pole double-throw switch. Naturally, other selection devices are within the level of ordinary skill in the art, including, for example, fuses, jumpers, DIP switches, and software-based selection devices.

A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made. For example, bi-directional bus 30 need not carry all communications between devices 10 and 20. For example, bi-directional bus 30 can include data lines but not control lines. Also, more than two devices can be connected to the bi-directional bus 30, and the devices can arbitrate for use of high and low frequency bands. In some implementations, more than two bands may be used.

Accordingly, other implementations are within the scope of the following claims.